

The blogosphere as an excitable social medium: Richter's and Omori's Law in media coverage

Peter Klimek^{a,b}, Werner Bayer^a, Stefan Thurner^{b,c,a,*}

^a*IIASA, Schlossplatz 1, A 2361 Laxenburg; Austria;*

^b*Section for Science of Complex Systems, Medical University of Vienna, Spitalgasse 23, A 1090 Vienna; Austria;*

^c*Santa Fe Institute; 1399 Hyde Park Road; Santa Fe; NM 87501; USA*

Abstract

We study the dynamics of public media attention by monitoring the content of online blogs. Social and media events can be traced by the propagation of word frequencies of related keywords. Media events are classified as exogenous – where blogging activity is triggered by an external news item – or endogenous where word frequencies build up within a blogging community without external influences. We show that word occurrences show statistical similarities to earthquakes. The size distribution of media events follows a Gutenberg-Richter law, the dynamics of media attention before and after the media event follows Omori's law. We present further empirical evidence that for media events of endogenous origin the overall public reception of the event is correlated with the behavior of word frequencies at the beginning of the event, and is to a certain degree predictable. These results may imply that the process of opinion formation in a human society might be related to effects known from excitable media.

1. Introduction

Modern man is exposed to a constant stream of news which we resorb, read, digest, discuss, disseminate and forget. Most news items are of relatively little impact. They receive a small amount of public attention over a short time and quickly descend into public oblivion. However, occasionally news reports have a massive impact so that they can overthrow public opinions, heralded beliefs and even governments. A recent example are the 2010-2011 Tunisian protests being sparked by reports of police use of tear gas against young demonstrators. Of course, the reasons for these protests go far beyond this single incident, they involve multiple political, social and economical dimensions. It is remarkable how reports about such relatively small incidents propagate under certain circumstances through a society – i.e. a strongly interconnected system – and trigger a nation-wide revolution, while under other circumstances the news will practically vanish unheard. It is tempting to see such media events as a human, social excitable medium. One may view them as a social analog to earthquakes [1, 2]. External stimuli trigger relaxation events where accumulated “energy” is discharged or spread within a complex, networked system. This phenomenon can also be observed in other excitable media such as the brain [3], oscillating chemical reactions like the Belousov-Zhabotinsky reaction [4] or the Mexican wave (or *La Ola*) at sport events [5].

Are there quantitative patterns in the way societies react to the arrival of breaking news? How can the impact

of a news item be quantified? How long and with which intensity do people devote their attention to current news? Such research questions become amenable to quantitative study through online news discussions. In recent years weblogging (or blogging) has emerged as a new publishing medium at the grassroots level of society. A blog is usually defined as a web page with entries listed in reverse chronological order, maintained by one or several writers. Blogs are often devoted to a unique topic, such as politics, finance or sports, and provide news or commentaries, often several times per day. In fall 2007 the blog search engine technorati.com stopped tracking the number of active blogs once it exceeded 100 million. This rapid development sparked academic interest. The dynamical evolution of the *blogosphere*, defined as the collection of all blogs at a given time, has been studied from a network perspective [6], where nodes are blogs which are connected if one blog possesses a hyper-link or URL to the other one. The spreading of news items can be seen as a diffusion process on this network [7]. The accelerating shift of human conversations and discussions toward online media like blogs, twitter, etc. has reached a point where web data mining techniques [8] can be used to e.g. predict spikes of sales ranks of books [9], measure public sentiments towards societal issues [10] or predict stock market movements [11].

In this work we study the impact of news reports on a small segment of the blogosphere by looking at timeseries of word frequencies. The sample contains blogs covering US domestic politics over a period of nearly two years. Blogs are typically not the first to report a story. They pick up developing news topics and offer comments and

*Corresponding author, stefan.thurner@meduniwien.ac.at

subjective points of view. They are therefore an ideal candidate to measure how the public ‘digests’ news. We show that the public reception of news reports follow a similar statistic as earthquakes do. The intensity of fore- and aftershocks can be described by a power law in analogy to Omori’s law [12], the size distribution of media events follows a Gutenberg-Richter law [13]. It has been reported previously that Omori’s law holds for the round-trip times of data packages in the internet [14]. Power law signatures in human activity have recently been found in the distributions of waiting times between a catastrophe and humanitarian responses [15] or between contributions to a discussion in online forums [16].

2. Data and Methods

A dataset of 168 large and massively popular blogs devoted to US domestic politics covering the entire political spectrum was compiled. The blog authors include popular political journalists (e.g. Glenn Beck or Taylor Marsh) as well as a number of self-pronounced political commentators labeling themselves everything from ‘far right’ to ‘liberal curmudgeons’. We recorded the content of each blog entry with a time stamp with a resolution of one second. We focus on the time period of 670 days between July 1st 2008 and May 3rd 2010. As the dominant themes this period contains the 2008 US presidential elections, the advent of Sarah Palin, the health care reform and the Iraq torture scandal. We developed a proprietary web crawler that continually crawls specified websites. The crawler can be targeted to a web page of a blog containing a blog entry of a given date. The content of this site is automatically parsed and checked for hyper-links to other blog entries. If a new link is found the linked page is downloaded and stored. The software then analyzes the structure of the blog entry, identifying and storing its date, headline, and entire content. Once all links in the web page containing the blog entry have been identified the initial blog page is searched for a link to another web page containing older entries of the same blog. This procedure is repeated as long as older entries are found. The crawler is implemented in the Java programming language in a fully object oriented fashion. All collected data is stored in a SQL database allowing efficient subsequent sorting and filtering. Filtering and pre-processing was done by removal of all short high frequency words e.g. ‘the’ or ‘and’, and removal of plural-s and other frequent endings such as ‘ing’ or ‘ed’.

We will refer to words surrounding political topics actively discussed in the blogosphere by ‘keywords’. They were extracted by the following procedure. The frequency of all $26^3 = 17576$ possible letter triplets (that is ‘aaa’, ‘aab’, ..., ‘zzz’) was counted over all blog entries for each day. Nearly 50% of these triplets were discarded since they never occurred. From the remaining triplets we extracted the time series of daily frequencies and looked at their co-occurrences with words. In particular, for each




Figure 1: (a): Word frequencies of ‘palin’ before and after the nomination of Sarah Palin as vice presidential candidate in the 2008 presidential elections. The x -axis is time (days) before and after the event, the event itself is indicated by a vertical dashed line. Note the absence of pre-cursory activity. (b): Word frequencies of ‘inauguration’ before and after the inauguration of Barack Obama. There is pronounced pre-cursory growth. Insets: same in log-log scale.

remaining triplet we searched for the day with the highest frequency, assume this day is t_0 . For this day a list of all words containing the triplet was made. For each word from this list the time series of daily word frequencies in the range between thirty days before and after t_0 was extracted. The word was kept for further analyses only if its daily word frequency has a maximum value within this range at t_0 . This procedure left us with approximately 4000 keywords.

Let $w_i(t)$ be the *word frequency* of keyword i at time t . An event is (somewhat arbitrarily) identified as a strong increase in word frequencies over a short period of time, with a subsequent decay to (usually) the previous levels. The time at which the peak occurs within an event is t_0 . The *event-size* E is defined by the peak level of word frequency relative to the average level within a timespan T before and after the peak level during the event. We work with $T = 30$. The event-size for word i at peak-time t_0 is

$$E_{i,t_0} = \frac{1}{2T+1} \frac{w_i(t_0)}{\sum_{t=t_0-T}^{t_0+T} w_i(t)} . \quad (1)$$

Two classes of response functions in social systems have been previously identified: endogenous and exogenous, see e.g. [17]. Endogenous events start with a phase of pre-cursory growth, exogenous ones by a sudden burst in word frequencies. Both types of events are followed by a relaxation process after the peak is reached. It usually follows a power law.

An example for an exogenous event is the nomination of Sarah Palin as vice presidential candidate in the 2008 US presidential elections. Fig. 1(a), shows the word frequencies of ‘palin’ before and after her nomination. The dashed line indicates the peak in word frequency, a sudden jump is followed by a relaxation process. As an example for an endogenous process consider the word frequencies of ‘inauguration’ before and after the inauguration of Barack Obama, Fig. 1(b). The day of the inauguration itself here coincides with $t_0 = 0$, the x -axis shows the days before




Figure 2: Cumulative distribution $\Pr(E^* > E)$ of event size E for exogenous and endogenous events. The exogenous case (blue circles, dotted line) can be fitted by a Gutenberg-Richter law with exponent $\beta_{exo} = 1.00(3)$ (green solid line). The exponent for the endogenous event sizes (red squares, dashed line) is $\beta_{exo} = 0.57(4)$ (black dash-dotted line), close to the typical value for earthquakes.

and after the inauguration. There is a phase of pre-cursory growth culminating at t_0 which is followed by a relaxation process.

It has been suggested [17] that word frequencies before and after an endogenous event at time t_0 follow power laws of the form

$$\begin{aligned} w_i(t < t_0) &\propto (t_0 - t)^{-\alpha_g} \\ w_i(t > t_0) &\propto (t - t_0)^{-\alpha_d} \end{aligned} \quad (2)$$

where in general the growth exponents before the peak at t_0 , α_g , and the decay exponent after t_0 , α_d , can be different. For the exogenous case there is no clear functional form for the word frequencies before t_0 , the relaxation dynamics often is of power law type [17],

$$w_i(t > t_0) \propto (t - t_0)^{-\gamma} \quad (3)$$

Exogenous events – translated into a seismological language – can be identified with the Omori law where $w_i(t > t_0)$ is interpreted as the analog to the aftershock rate. The idea is that the frequencies of mentionings of a keyword before (after) t_0 can be regarded as foreshock (aftershock) rates of that particular event. For endogenous events we also observe ‘foreshocks’ following an Omori law.

For each event in a timeseries of word frequencies in our database it was checked whether it is endogenous or exogenous. For this we segmented each timeseries (often containing several events) into windows of T days and located the local maxima within this timespan¹. The event was classified as endogenous or exogenous if the Residual Sum of Squares for the fit was below 0.15, computed by fits to Eqs. (2) and (3) around t_0 , and normalized by $w_i(t_0)$.

¹To fit a power growth or decay we required at least non-zero word frequencies over two weeks before and after t_0 . The range of the fit was chosen between $(t_0, t_0 \pm \tau)$ with $\tau \in \{14, \dots, 30\}$. For each value of τ in this range the Akaike Information Criterion was computed. The exponent of the fit with minimal value of this criterion was used for further analyses.




Figure 3: Cumulative distribution functions of γ (blue dotted line), α_g (red solid) and α_d (green dashed). The similar curves for γ and α_d suggest a universal decay law of public attention for media events. The exponents for pre-cursory growth α_g are typically higher than the decay exponents α_d .

3. Results

With the above procedure we ended up with approximately 150 endogenous and 1000 exogenous events, i.e. an average of 0.2 endogenous and 1.5 exogenous events per day. Note that exogenous events are about an order of magnitude more frequent. One media event generally corresponds to events in several word frequencies of related keywords. This effect should be the same for endogenous and exogenous types.

The Gutenberg-Richter law is an empirical power law describing the frequency of earthquakes of a given radiated seismic energy E_s ,

$$\Pr(E_s^* > E_s) \propto E_s^{-(\beta+1)} \quad (4)$$

with a typical exponent of $\beta \approx 2/3$ [12]. Figure 2 shows the cumulative distribution function (CDF) $\Pr(E^* > E)$ for endogenous and exogenous events. In both cases we find a Gutenberg-Richter law

$$\Pr(E^* > E) \propto E^{-(\beta+1)} \quad (5)$$

For endogenous events we find $\beta_{endo} = 0.57(4)$, for exogenous $\beta_{exo} = 1.00(3)$. Together with the relative high number of exogenous compared to endogenous events, this suggests that exogenous events share characteristics of ‘earthquake swarms’ [18] – sequences of small earthquakes over a relatively short time period.

The cumulative distribution functions for γ for exogenous events and α_d and α_g for endogenous events are shown in Fig. 3. The value of the CDF of γ^* at any given value γ is defined as the probability that γ^* is greater than γ , that is $\Pr(\gamma^* \geq \gamma)$. Exponents γ and α_d follow almost the same distribution function. This might suggest a universal decay law of public attention to media events. For high values of exponents the CDF of growth exponents α_g is larger than the distribution function for decay exponents. In general we find for a given exponent value χ

$$\Pr(\chi > \alpha_g) > \Pr(\chi > \alpha_d) \approx \Pr(\chi > \gamma) \quad (6)$$




Figure 4: Scatter plot of α_g versus α_d . Each point is one endogenous event with x - and y -coordinates given by the values of the α 's in Eqns. 2. The dashed line is a regression line, we find a correlation coefficient of $\rho \approx 0.67$ and a p -value of 10^{-22} against the null hypothesis of no correlation.

To what extent can one predict the dynamics of endogenous events from their pre-cursory growth? Fig.4 shows for each endogenous event its growth α_g versus its decay exponent α_d . Especially for low values of the exponents (smaller than e.g. one) there is a clear correlation between them. Even if we include the entire dataset of endogenous events we can reject the null hypothesis of no correlation up to a p -value of 10^{-22} . The correlation coefficient is $\rho \approx 0.67$. This correlation should drastically increase if one would only include events with small growth and/or decay exponents α .

4. Conclusions

Blogs offer a new and exciting possibility to study collective human behavior on a quantitative basis. The blogosphere is a highly connected virtual space where people with different background and attitude disseminate and discuss issues that caught their attention and sufficient interest, in a way that can be exploited for quantitative studies. We empirically studied how people react to new pieces of information through dynamical patterns of their blogging behavior. We see the blogosphere is a human, excitable social medium in which waves of collective excitement can be traced. We studied event-size, foreshock and aftershock distributions in this medium and noticed analogies to seismology. The intensity of fore- and aftershocks follows Omori's law, the distribution of event-sizes is of Gutenberg-Richter type. We presented indications that there exist significant correlations between the dynamics of fore- and aftershocks. One might also think of a 'Richter scale' for media events. The largest event recorded in our dataset is the nomination of Sarah Palin as vice presidential candidate. Indeed, aftershocks of this event are still trembling and quivering through our society.

References

- [1] B. Gutenberg, C.F. Richter, Frequency of earthquakes in California, *Bull. Seismol. Soc. Am.* 34 (1944), 4, 185-188.
- [2] Y.Y. Kagan, D.D. Jackson, Long-term earthquake clustering, *Geophys. J. Int.* 104 (1991), 1, 117-134.
- [3] I. Osorio, M.G. Frei, D. Sornette, J. Milton, Y.C. Lai, Epileptic seizures: Quakes of the brain? *Phys. Rev. E*, 82 (2010), 2, 021919.
- [4] B. P. Belousov A periodic reaction and its mechanism, *Compilation of Abstracts on Radiation Medicine*, 147 (1959), 145.
- [5] I. Farkas, D. Helbing, T. Vicsek, Mexican waves in an excitable medium, *Nature*, 419 (2002), 131-132.
- [6] R. Kumar, J. Novak, P. Raghavan, A. Tomkins, On the Bursty Evolution of Blogspace, *World Wide Web*, 8 (2005), 2, 159-178.
- [7] D. Gruhl, R. Guha, D. Liben-Nowell, A. Tomkins, Information Diffusion Through Blogspace, In *Proceedings of the 13th international conference on World Wide Web* (2004), 491-501, NY, USA.
- [8] B. Liu, *Web Data Mining*, first ed., Berlin, Springer, 2006.
- [9] D. Gruhl, R. Guha, R. Kumar, J. Novak, A. Tomkins, The Predictive Power of Online Chatter, *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining* (2005), 78-87, Chicago, Illinois, USA.
- [10] Y. Liu, X. Huang, A. An, X. Yu, ARSA: A Sentiment-Aware Model for Predicting Sales Performance Using Blogs, *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval* (2007), 607-614, Amsterdam, Netherlands.
- [11] W. Zhang, S. Skiena, Trading Strategies To Exploit Blog and News Sentiment, *The 4th International AAAI Conference on Weblogs and Social Media* (2010), Washington, DC.
- [12] T. Utsu, Y. Ogata, S. Matsu'ura, The Centenary of the Omori Formula for a Decay Law of Aftershock Activity, *J. Phys. Earth*, 43 (1995), 1-33.
- [13] L. Knopoff, Y.Y. Kagan, R. Knopoff, b values for foreshocks and aftershocks in real and simulated earthquake sequences. *Bull. Seism. Soc. Am.*, 72 (1982), 1663-1676.
- [14] S. Abe, N. Suzuki, Omori's law in the Internet traffic, *Europhys. Lett.*, 61 (2003), 852-855.
- [15] R. Crane, F. Schweitzer, D. Sornette, Power law signature of media exposure in human response waiting time distributions, *Phys. Rev. E*, 81 (2010), 5, 056101.
- [16] M. Mitrovic, B. Tadic, Bloggers behavior and emergent communities in Blog space, *Eur. Phys. J. B*, 73, (2010), 293-301.
- [17] R. Crane, D. Sornette, Robust dynamic classes revealed by measuring the response function of a social system, *Proc. Nat. Acad. Sci.*, 105 (2008), 41, 15649-15653.
- [18] L.R. Sykes, Earthquake swarms and sea-floor spreading, *J. Geophys. Res.*, 75 (1970), 6598-6611.